Automatically Reinforcing a Game AI

David L. St-Pierre, Jean-Baptiste Hoock, Jialin Liu, Fabien Teytaud and Olivier Teytaud

arXiv:1607.08100v1 [cs.AI] 27 Jul 2016

***Abstract*—A recent research trend in Artificial Intelligence (AI) is the combination of several programs into one single, stronger, program; this is termed portfolio methods. We here investigate the application of such methods to Game Playing Programs (GPPs). In addition, we consider the case in which only one GPP is available - by decomposing this single GPP into several ones through the use of parameters or even simply random seeds.**

These portfolio methods are trained in a learning phase. We propose two different offline approaches. The simplest one, BestArm, is a straightforward optimization of seeds or parameters; it performs quite well against the original GPP, but performs poorly against an opponent which repeats games and learns. The second one, namely Nash-portfolio, performs similarly in a "one game" test, and is much more robust against an opponent who learns. We also propose an online learning portfolio, which tests several of the GPP repeatedly and progressively switches to the best one - using a bandit algorithm.

***Index Terms*—Monte Carlo Search, Nash Equilibrium, Portfolios of policies.**

I. Introduction

Portfolios are widely used in many domains; after early papers in machine learning [1], [2], they are now ubiquitous in Artificial Intelligence, planning, and combinatorial optimization [3]–[5]. The special case of parameter tuning (close to our "variants problem" later in the present document) is widely studied [6], with applications to SAT-solving [7], [8] or computer vision [9].

Recently, portfolios were also applied in games [10], [11].

A "portfolio" here refers to a family of algorithms which are candidates for solving a given task. On the other hand, "portfolio combination" or "combination" refers to the combined algorithm. Let us introduce a simple combined algorithm. If we have algorithms $\pi_1, \ldots, \pi_K$ in the portfolio, and if the combination is $\pi = \pi_i$ with probability p_i where $0 \leq p_i \leq 1$ and $\sum_{i=1}^{K} p_i = 1$ (the random choice is made once and for all at the beginning of each game), then π is, by definition, the portfolio combination with probability distribution p. Moreover, also by definition, it is stationary. Furthermore we will consider a case in which the probability distribution is not stationary (namely, UCBT, defined in Section III-B).

Another approach, common in optimization, is "chaining" [12], which means interrupting one program and using its internal state as a hint for another algorithm. The combination can even be "internal" [13], i.e. parts of a solver are used in other solvers. The most famous applications of portfolios are in SAT-solving [14]; nowadays, portfolios routinely win SAT-solving competitions.

In this paper, we focus on portfolios of policies in games, i.e. portfolios of GPP. Compared to optimization, portfolios of policies in games or control policies have been less widely explored, except for e.g. combinations of local controllers by Fuzzy Systems [15], Voronoi controllers [16] or some case-based reasoning [17]. These methods are based on "internal" combinations, using the current state for choosing between several policies. We here focus on external combinations; one of the internal programs is chosen at the beginning of a game, for all games. Such combinations are sometimes termed "ensemble methods"; however, we simply consider probabilistic combinations of existing policies, the simplest case of ensemble methods. This is an extension of a preliminary work [18].

To the best of our knowledge, there is not much literature on combining policies for games when only one program is available. The closest past work might be Gaudel et al. [19], which proposed a combination of opening books, using tools similar to those we propose in Section III-A for combining policies.

A. Main goal of the present paper

The main contribution of this paper is to propose a methodology that can generically improve the performance of policies without actually changing the policies themselves, except through the policy's options or the policy's random seed. Incidentally, we establish that the random seed can have a significant contribution to the strength of an artificial intelligence, just because random seeds can decide the answer to some critical moves as soon as the original randomized GPP has a significant probability of finding the right move. In addition, while a fixed random seed cannot be strong against an adaptive opponent, our policies are more diversified (see the Nash approach) or adaptive (see our UCBT-portfolio).

Our approach is particularly relevant when the computational power is limited, because the computational overhead is very limited. Our main goal is to answer the following question: how can we, without development and without increasing the online computational cost, significantly increase the performance of a GPP in games ?

B. Outline of the present paper

We study 2 different portfolio problems:

- The first test case is composed of a set of random seeds for a given GPP. By considering many possible seeds, we get deterministic variants of the original stochastic

D.L. St-Pierre is with the Department of Industrial Engineering, Univ. du Québec à Trois-Rivières, Trois-Rivières, Qc, G9A 5H7 CAN. E-mail: lupienst@uqtr.ca.

J.-B. Hoock is with TAO (Inria), LRI, Univ. Paris-Sud, Paris, France. E-mail: jbhoock@gmail.com.

J. Liu is with the School of Computer Science and Electronic Engineering, Univ. of Essex, Wivenhoe Park, CO4 3SQ, UK. E-mail: jialin.liu@essex.ac.uk.

F. Teytaud is with Univ. Lille Nord de France, ULCO, LISIC, Calais, France. E-mail: teytaud@lisic.univ-littoral.fr.

O. Teytaud is with Google Zürich, Brandschenkestrasse 110, 8002 Zrich, Switzerland. E-mail: olivier.teytaud@gmail.com.

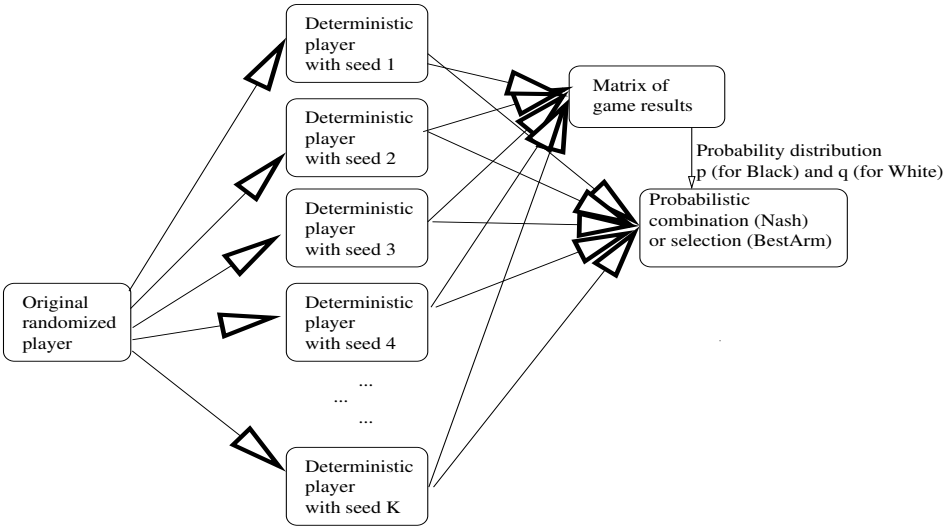
Deterministic player with seed 1
Deterministic player with seed 2
Deterministic player with seed 3
Deterministic player with seed 4
...
Deterministic player with seed K
Original randomized player
Matrix of game results
Probability distribution p (for Black) and q (for White)
Probabilistic combination (Nash) or selection (BestArm)

Fig. 1: Method used for generating a portfolio of deterministic programs from a randomized one (left part of the figure) and combining them back into one single randomized program better than its original self. The UCBT portfolio proposed in the present paper does not directly fit in this figure because it depends on earlier games: it is non stationary.

GPP. We restrict our attention to combinations which are a fixed probability distribution over the portfolio: we propose a combination such that, at the beginning of the game, one of the deterministic GPPs (equivalently, one of the seeds) is randomly drawn and then blindly applied. Hence, the problem boils down to finding a probability distribution over the set of random seeds such that it provides a strong strategy. We test the obtained probability distribution on seeds versus (i) the GPP with uniformly randomly drawn seeds (i.e. the standard, original, version of the GPP) and (ii) a stronger GPP, defined later, termed "exploiter" (Section V-A1).

- In the second case, we focus on different parameterizations of a same program, so that we keep the spirit of the main goal above. The goal here is to find a probability distribution over these parameterizations. We will assess the performance of the obtained probability distribution against the different options.

A combination can be constructed either offline [20] or online [21], [22]. In this paper, we use three different methods for combining several policies:

- In the first one, termed Nash-portfolio, we compute a Nash Equilibrium (NE) over the portfolio of policies in an offline fashion. This approach computes a distribution such that it generates a robust (not exploitable) agent. Further tests show a generalization ability for this method.
- In the second one, termed UCBT-portfolio, we choose an element in the portfolio, online, using a bandit approach. This portfolio learns a specialized distribution, adaptively, given a stationary opponent. This approach is very good at exploiting such opponent.
- The third one, $Best\ Arm$, is the limit case of UCBT-portfolio. It somehow cheats by selecting the best option against its opponent, i.e. it uses prior knowledge. This is what UCBT will do asymptotically, if it is allowed to play enough games.

These concepts are explained in Fig. 1. There are important related works using teams of programs [23]–[25]. The specificity of the present work is to get an improvement with a portfolio of programs which are indeed obtained from a single original program - i.e. we get an improvement "for free", in terms of development.

The rest of the paper is divided as follows. Section II formalizes the problem. Section III describes our approach. Section IV details the experimental setup. Section V presents the results. Section VI shows robustness elements. Section VII presents simplified variants of our algorithms, performing similarly to the original ones. Section VIII concludes.

II. Problem Statement

In this section, we formalize the notion of policies, adversarial portfolios, and the framework of matrix games. We also introduce the concepts of overfitting, exploitation and generalization.

A. Policies

We consider policies, i.e. game playing programs (GPP [24]), and tools (portfolios) for combining/selecting them.

When a GPP is stochastic, it can be made deterministic by choosing a fixed seed at the beginning of the game. From a stochastic π, we can therefore build several GPP π_1, π_2, ... corresponding to seeds 1, 2, ... In the case of our portfolio, and in all experiments and algorithms in the present paper, the choice of the seed is done once and for all, when a new game starts.

B. Matrix games

In this paper we only consider finite constant-sum adversarial games (i.e. if one player wins the other loses, constant-sum and adversarial are synonyms) with a reward that is only available at the end of the game. To properly define our algorithms in the following sections, let us introduce the concept of constant-sum matrix game. Without loss of generality, we define the concept of 1-sum matrix game instead of an arbitrary constant.

Consider a matrix $K \times K'$, with values in $[0, 1]$. This matrix models a game as follows:

- Simultaneously and privately:
  - Player 1 chooses $i \in \{1, \dots, K\}$.
  - Player 2 chooses $j \in \{1, \dots, K'\}$.
- Then they receive rewards as follows:
  - Player 1 receives reward $M_{i,j}$.
  - Player 2 receives reward $1 - M_{i,j}$.

A pure strategy (for player 1) consists in playing a given, fixed $i \in \{1, \dots, K\}$, with probability 1. A mixed strategy, or simply a strategy, consists in playing i with probability p_i, where $\sum_{i=1}^{K} p_i = 1$ and $\forall i \in \{1, \dots, K\}, 1 \geq p_i \geq 0$. Pure and mixed strategies for player 2 are defined similarly. Pure strategies are a special case of mixed strategies.

In the general stationary case, Player 1 chooses row i with probability p_i and Player 2 chooses column j with probability q_j. It is known since [26], [27] that there exist strategies p and q for the first and second player respectively, such that

$$\forall (p', q'), p'Mq \leq pMq \leq pMq'. \tag{1}$$

p and q are not necessarily unique, but the value $v = pMq$ is unique (this is a classical fact, which can be derived from Eq. 1) and it is, by definition, the value of the game. The exploitability of a strategy p' for the first player is $exploit_1(p') = v - \min_q p'Mq$. When $exploit_1(p') = 0$, it is equivalent to the fact that $p' = p$. The exploitability of a strategy q' for the second player is $exploit_2(q') = \max_p pMq' - v$ and it verifies similar properties. The exploitability of a strategy is always non-negative and quantifies the robustness of a strategy. The exploitability of a GPP which can play both as Player 1 and as Player 2 is the average of its exploitabilities as Player 1 and its exploitability as Player 2.

C. Overfitting, exploitation & generalization

Overfitting in a game sense refers to the poor performance of a GPP P when P seems to be strong according to a given criterion which was used in the design of P. For instance, a GPP built through trials and errors by accepting any modifications which increase the success rate against a GPP X might have an excellent success rate against X, but a poor winning rate against another program Y. This is a case of overfitting.

This is important when automatic tuning is applied, and in particular for portfolio methods when working on random seeds. Selecting good random seeds for Player 1, by analyzing a matrix of results for various seeds for Player 1 and Player 2, might be excellent in terms of performance against the seeds used for Player 2 in the data; but for a proper assessment of the performance against the original randomized program, we should use games played against other seeds for Player 2. The performance against the seeds used in the data is referred to as an empirical performance, whereas the performance against new seeds is referred to as the performance in generalization [28]. Only the performance in generalization is a proper assessment of performance; we provide such results.

In games, overfitting is related to exploitability. Exploitability is an indicator of overfitting; when we build a GPP by some machine learning method, we can check, by the exploitability measure, whether it is good more generally than just against the opponents which have been used during the learning process.

In practice, exploitability defined as above is hard to measure. Therefore, we often use simpler proxies, e.g. the worst performance against a set of opponents. We say that a program A "exploits" a program B when A has a great success rate against B, much higher than the success rate of most programs against B - and we say that a family $\mathcal{A}$ exploits a program B when there exists $A \in \mathcal{A}$ which exploits B. The existence of A which "exploits" B suggests an overfitting issue in the design of B.

III. Approaches

Section III-A proposes a method for combining policies offline, given a set of policies for Player 1 and a set of policies for Player 2. Section III-B proposes a method for combining policies online, given a portfolio of policies for player 1 and a stationary opponent.

A. Offline learning: Nash-portfolios and $Best\ Arm$

Consider two players P_1 and P_2, playing some game (not necessarily a matrix game). P_1 is Black, P_2 is White. Assume that P_1 has a portfolio of K policies. Assume that P_2 has a portfolio of K' policies. Then, we can construct a static combination of these policies by solving (i.e. finding a Nash equilibrium of) the matrix game associated to the matrix M, with $M_{i,j}$ the winning rate of the i^{th} policy of P_1 against the j^{th} policy of P_2. Solving this 1-sum matrix game provides $p_1, \dots, p_K$ and $q_1, \dots, q_{K'}$, probabilities, and the combination consists in playing, for P_1, the i^{th} policy with probability p_i and, for P_2, the j^{th} policy with probability q_j. Such a combination will be termed here a Nash-portfolio. By construction,

- the Nash-portfolio can play both as Black and as White (P_1 and P_2);
- the Nash-portfolio does not change over time but is, in the general case, stochastic.

Let us define more formally the Nash-portfolio and the $Best\ Arm$ portfolio.

Definition: *Given a set S_1 of K policies for Black and a set S_2 of K' policies for White. Define $M_{i,j}$ the winning rate of the i^{th} strategy in S_1 against the j^{th} strategy in S_2. Then the strategy which plays:*

- *as Black, the i^{th} strategy in S_1 with probability p_i;*

- *as White, the j^{th} strategy in S_2 with probability q_j;*

is termed a $Nash$-portfolio of (S_1, S_2) if (p, q) is a solution of Eq. 1.

The strategy playing the I^{th} strategy in S_1 with probability 1 when playing Black, and playing the J^{th} strategy in S_2 with probability 1 when playing White, is a $Best\ Arm$ portfolio if I maximizes

$$\sum_{j=1}^{K'} M_{I,j} \tag{2}$$

and J minimizes

$$\sum_{i=1}^{K} M_{i,J}. \tag{3}$$

The strategy playing the i^{th} strategy in S_1 as Black (resp. in S_2 as White) with probability $1/K$ (resp. $1/K'$) is the uniform portfolio.

$Best\ Arm$ can be seen as the best response to the uniform policy. In both cases, Nash-portfolio and $Best\ Arm$, there is no uniqueness.

The Nash equilibrium can be found using an exact solving, in polynomial time, by linear programming [29]. It can also be found approximately and iteratively, in sublinear time, as shown by [30], [31]; the EXP3 algorithm is classical for doing so.

From the properties of Nash equilibria, we deduce that the Nash-portfolio has the following properties:

- It depends on a family of policies for player 1 *and* on a family of policies for player 2. It is therefore based on a training, by offline learning.
- It is optimal (for player 1) among all mixed strategies (i.e. stochastic combinations of policies in the portfolio of player 1), in terms of both
  - worst case among the pure strategies in the portfolio of player 2;
  - worst case among the mixed strategies over the portfolio of player 2.
- It is not necessarily uniquely defined.

In optimization settings, it is known [32] that having a somehow "orthogonal" portfolio of algorithms, i.e. algorithms as different from each other as possible, is a good solution for making the combination efficient. It is however difficult, in the context of policies, to know in advance if two algorithms are orthogonal - we can however see, a posteriori, which strategies have positive probabilities in the obtained combination.

B. Online learning: UCBT-Portfolio

Section III-A assumed that S_1 and S_2, two sets of strategies, are available and that we want to define a combination of policies in S_1 (resp. in S_2). A different point of view consists in adapting online the probabilities p_i and q_i, against a fixed opponent. We propose the following algorithm. We define this approach in the case of Black, having K policies at hand. The approach is similar for White. It is directly inspired by the bandit literature [33], [34], and, more precisely, by Upper-Confidence-Bounds-Tuned (UCBT) [35], with parameters optimized for our problem:

- Define $n_i = 0$, $r_i = 0$, for $i \in \{1, \dots, K\}$.
- For each iteration $t \in \{1, 2, 3, \dots\}$.
  - compute for each $i \in \{1, \dots, K\}$ $score(i) = \min(1, r_i/n_i + \frac{1}{100}\sqrt{C \log(4t^p)/n_i}$ $+\frac{16}{100}\log(4t^p)/n_i)$. using $X/0 = +\infty$ (even for $X = 0$), $p = 2.1$ and $C = 2$ (UCBT, i.e. UCB-Tuned, formula).
  - choose k maximizing $score(k)$.
  - play a game using algorithm k in the portfolio.
  - if it is a win, $r_k \leftarrow r_k + 1$.
  - $n_k \leftarrow n_k + 1$.

Definition. *We refer to this adaptive player as UCBT-Portfolio, or Bandit-Portfolio.*

IV. Settings

This section presents the settings used in our experiments. Section IV-A details the notion of portfolio of random seeds for 4 different games (Go, Chess, Havannah, Batoo). Section IV-B explains the context a portfolio of parameterizations for the game of Go.

A. Portfolio of Random Seeds

First, let us explain the principle of GPPs that just differ by their random seeds. We first apply the portfolio approach in this case. Without loss of generality, we will focus on the case where $K = K'$. The K GPPs for Black and the K GPPs for White use random seed 1, 2, ..., K respectively. Let us see what our $Nash$-portfolio and other portfolios become in such a setting. We define $M_{i,j} = 1$ if, with random seed i, Black wins against White with random seed j. Otherwise, $M_{i,j} = 0$. Importantly, the number of games to be played for getting this matrix M, necessary for learning the Nash-portfolio is K^2. This is because there is no need for playing multiple games, since fixing the random seed makes the result deterministic. Thus, we just play one game for each $(i, j) \in \{1, \dots, K\}^2$.

Then, we compute (p, q), one of the Nash equilibria of the matrix game M. This learns simultaneously the Nash-portfolio for Black and for White. Using this matrix M, we can also apply:

- the uniform portfolio, simply choosing randomly uniformly among the seeds;
- the $Best\ Arm$ portfolio, choosing (I, J) optimizing Eqs. 2 and 3 and using I as a seed for Black and J as a seed for White;
- the UCBT-portfolio, which is the only non-stationary portfolio in the present paper.

We use 4 different testbeds in this category (portfolio of random seeds): Go, Chess, Havannah, Batoo. These games are all deterministic (Batoo has an initial important simultaneous move, namely the choice of a base-build, i.e. some initial stones - but we do not keep the partially observable stone, see details below).

1) The game of Go: The first testbed is the game of Go for which the best programs are Monte-Carlo Tree Search (MCTS) with specialized Monte Carlo simulations and patterns in the tree. The Black player starts. The game of Go is an ancient

oriental game, invented in China probably at least 2 500 years ago. It is still a challenge for GPP, as even though MCTS [36] revolutionized the domain, the best programs are still not at the professional level. Go is known as a very deep game [37].For the purpose of our analysis, we use a 9x9 Go board.

We use GnuGo's random seed for having several GnuGo variants. The random seed of GnuGo makes the program deterministic, by fixing the seed used in all random parts of the algorithm. We define 32 variants, using "GnuGo –level 10 –random-seed k" with $k \in \{1, \ldots, 32\}$. In other words, we use a MCTS with 80 000 simulations per move, as GnuGo uses, by default, 8 000 simulations per level.

2) Chess: The second testbed is Chess. There are 2 players, Black and White. The White player starts. Chess is a two-player strategy board game played on a chessboard, a checkered game board with 64 squares arranged in an 8-by-8 grid. As in Go, this game is deterministic and full information. For the game of Chess, the main algorithm is alpha-beta [38], yet here we use a vanilla MCTS. We define 100 variants for the portfolios of random seeds (giving a matrix M of size 100-by-100), using a MCTS with 1 000 simulations per move and enhanced by an evaluation function. Our implementation is roughly ELO 1600 on game servers, i.e. amateur level.

3) Havannah: The third testbed is the game of Havannah. There are 2 players in this game: Black and White. The Black player starts. Havannah is an abstract board game invented by Christian Freeling. It is best played on a base-10 hexagonal board, i.e. 10 hexes (cells) to a side. Havannah belongs to the family of games commonly called connection games; its relatives include Hex and TwixT. This game is also deterministic with full information. For the game of Havannah, a vanilla MCTS with rapid action value estimates [39] provides excellent performance. We define 100 variants for the portfolio of random seeds (giving a matrix M of size 100-by-100), using a MCTS with 1 000 simulations per move.

4) Batoo: The fourth testbed is a simplified version of Batoo. Batoo is related to the game of Go, but contains 2 features which are not fully observable:

- Each player, once per game, can put a hidden stone instead of a standard stone.
- At the beginning, each player, simultaneously and privately, puts a given number of stones on the board. These stones, termed "base build", define the initial position. When the game starts, these stones are revealed to the opponent and colliding stones are removed.

We consider a simplified Batoo, without the hidden stones - but we keep the initial, simultaneous, choice of base build. As in Go, this game is deterministic. Once the initial position of the stones is chosen for both player a normal game of 9x9 Go is executed using a GnuGo level 10.

B. Portfolio of parameterizations: variants of GnuGo

We consider the problem of combining several variants (each variant corresponds to a set of options which are enabled) of a GPP for the game of Go.

Our matrix M is a 32×32 matrix, where $M_{i,j}$ is the winning rate of the i^{th} variant of GnuGo (as black) against the j^{th} variant of GnuGo (as white). We consider all combinations of 5 options of GnuGo, hence $32 = 2^5$ variants. In short, the first option is 'cosmic-go', which focuses on playing at the center. The second option is the use of fuseki (global opening book). The third option is 'mirror', which consists in mirroring your opponent at the early stages of the game. The fourth option is the large scale attack, which evaluates if a large attack across several groups is possible. The fifth option is the break-in. It consists in breaking the game analysis into territories that require deeper tactical reading and are impossible to read otherwise. It revises the territory valuations. Further details on the 5 options are listed on our website [40].

As opposed to Section IV-A, we need more than one evaluation in order to get $M_{i,j}$, because the outcome of a game between 2 different GPP is not deterministic. For the purpose of this paper, we build the matrix $M_{i,j}$ offline by repeating each game (i, j) 289 times, leading to a standard deviation at most 0.03 per entry.

For this part, experiments are performed on the convenient 7x7 framework, with MCTS having 300 simulations per move - this setting is consistent with the mobile devices setting. We refer to the i^{th} algorithm for Black as BAIi (Black Artificial Intelligence # i), and WAIj is the j^{th} algorithm for White.

V. Experiments

In this section we evaluate the performance of our approaches across different settings.

Section V-A focuses on the problem of computing a probability distribution in an offline manner for the games defined in Section IV. We evaluate the scores of the Nash-portfolio approach and of the $Best\ Arm$ approach. We also include the uniform portfolio. In the case of a portfolio of random seeds, the uniform portfolio is indeed the original algorithm.

Section V-B focuses on the problem of learning a probability distribution in an online manner to play against a specific opponent for the games defined in Section IV. We evaluate the learning ability of our UCBT-portfolio.

A. Learning Offline

In this section we present an analysis of the different offline portfolios across the testbeds. Table I shows the performance of the portfolios. The column V presents the value of the matrix game M. The following columns are self-explanatory where 1 indicates the player with the initiative and 2 indicates the player without.

We briefly describe the results in the four portfolios of random seeds as follows:

- For the game of Go, the number of seeds with positive probability in the Nash-portfolio is 11 for Black and 9 for White, i.e. roughly $\frac{1}{3}$ of the random seeds.Nash-portfolio outperforms $Best\ Arm$, which in turn wins against Uniform.
- In Chess, the number of seeds with positive probability in the Nash equilibrium is 34 for White and 37 for Black, i.e. roughly $\frac{1}{3}$ of the random seeds. The best arm strategy is easily beaten by the Nash portfolio.

TABLE I: Performance Portfolio Analysis. The Nash-portfolio clearly outperforms the uniform one (which is the original algorithm), but not necessarily the simple BestArm algorithm; BestArm has some weaknesses in terms of exploitability (as discussed later, in Fig. 2(b)) but it is not necessarily weaker than Nash for direct games one against each other.

	V	Nash(1) vs Unif(2)	Nash(2) vs Unif(1)	Nash(2) vs $Best\ Arm(1)$	Nash(1) vs $Best\ Arm(2)$
Go	54.16%	68.51%	38.80%	55.76%	66.62%
Chess	54.52%	59.16%	50.29%	80.31%	86.07%
Havannah	55.36%	58.10%	52.51%	72.69%	75.75%
Batoo	50.11%	70.69%	34.01%	56.56%	67.95%
Variants	61.21%	65.57%	52.37%	61.21%	70.52%

- For the game of Havannah, the number of seeds with positive probability in the Nash-portfolio is 36 for White and 34 for Black, i.e. roughly $\frac{1}{3}$ of the random seeds are selected. The best arm strategy is outperformed by the Nash portfolio.
- For the game of Batoo, the number of seeds with positive probability in the Nash-portfolio is 11 for Black and 14 for White. The uniform strategy is seemingly quite easily beaten by the Nash-portfolio or the $Best\ Arm$-portfolio.

These descriptive statistics are extracted in the learning step, i.e. on the training data, namely the matrix M. They provide insights, in particular around the fact that no seed dominates, but a clean validation requires a test in generalization, as discussed in Section II-C. Performances in generalization are discussed in Section V-A1.

We now consider the case of "Variants", which refers to the case in which we do not work on random seeds, but on variants of GnuGo, as explained in Section IV-B; the goal is to "combine" optimally the variants, among randomized choices between variants. For Variants, in the NE, the number of selected options (i.e. options with positive probability) is 4 for Black and also 4 for White, i.e. $\frac{1}{8}$ of the variants are in the Nash. This means that no option could dominate all others. The uniform strategy is quite easily beaten by the Nash as well as the best arm strategy. The $Best\ Arm$ portfolio is beaten by the $Nash$ portfolio. This last point is interesting: it shows that selecting the variant which is the best for the average winning rate against other variants (this is what $Best\ Arm$ does), leads to a combined variant which is weaker than the Nash combination.

1) Generalization ability of offline Portfolio: We now switch to the performance in generalization of the Nash and $Best\ Arm$ approach. In other words, we test whether it is possible to use a distribution computed over a portfolio of policies (learned against a given set of opponent policies) against new opponent policies that are not part of the initial matrix. The idea is to select a submatrix of size K (learning set), compute our probability distribution for this submatrix using either $Nash$ or $Best\ Arm$ and make it play against the remainder of the seeds (validation set). We restrict our analysis to the setting presented in Section IV-A. We focus on the 4 portfolios with random seeds. We test policies (Nash-portfolio, uniform portfolio, $Best\ Arm$) against an opponent which is not in the training set in order to evaluate whether our approach is robust.

The x-axis represents the number of policies K considered for each player (hence a matrix M of type $K \times K$). The y-axis shows the win rate of the different approaches

- against an opponent that uses the uniform strategy (this is tested with independently drawn random seeds, not used in the matrix used for learning);
- against an "exploiter"; by exploiter, we mean an algorithm which selects, among the $N > K$ considered seeds which are not used in the learning set, the best performing one. Obviously, this opponent is somehow cheating; he knows which probability distribution you have, and uses it for choosing his seed among the $M = N - K$ seeds which are considered in the experiment but not used in the learning. This is a proxy for the robustness; some algorithms are better than other for resisting to such opponents who use some knowledge about you.

Figure 2 summarizes the results for the game of Go. Figure 2(a) presents the results of 2 different approaches (Nash and $Best\ Arm$) versus the uniform baseline. All experiments are reproduced 10 000 times (5 000 times for the Black player and 5 000 times for the White player) and standard deviations are smaller than 0.007.

Figure 2(b) shows the difference between a $Nash$ approach and a $Best\ Arm$ approach in terms of exploitability.

From Figure 2(a) we can observe that there is a clear advantage to use either the $Nash$ or the $Best\ Arm$ approach when facing a new set of policies. Moreover, as expected, as the size of the initial matrix grows, the winning rates of both $Nash$ and $Best\ Arm$ increase when compared to the baseline. It is interesting to note that there is a sharp increase when the submatrix size is relatively small (between 3 and 7). Afterwards, the size of the submatrix has a moderate impact on the performance until most options are included in the matrix.

It does not come as a surprise that the approach $Best\ Arm$ performs slightly better than the $Nash$ against a uniformly random opponent. The $Best\ Arm$ approach is particularly well suited to play against such an opponent. However, the $Best\ Arm$ approach is easily exploitable. This behavior is shown in Figure 2(b).

From Figure 2(b) it clearly appears that $Best\ Arm$ is a strategy very easy to exploit. Thus, even if Figure 2(a) shows that the use of the $Best\ Arm$ approach outperforms $Nash$ versus the uniform baseline, $Nash$ is a much more resilient strategy.

Chess: Figure 3 summarizes the results for the game of Chess. Figure 3(a) presents the results of 2 different approaches versus the uniform baseline. All experiments are

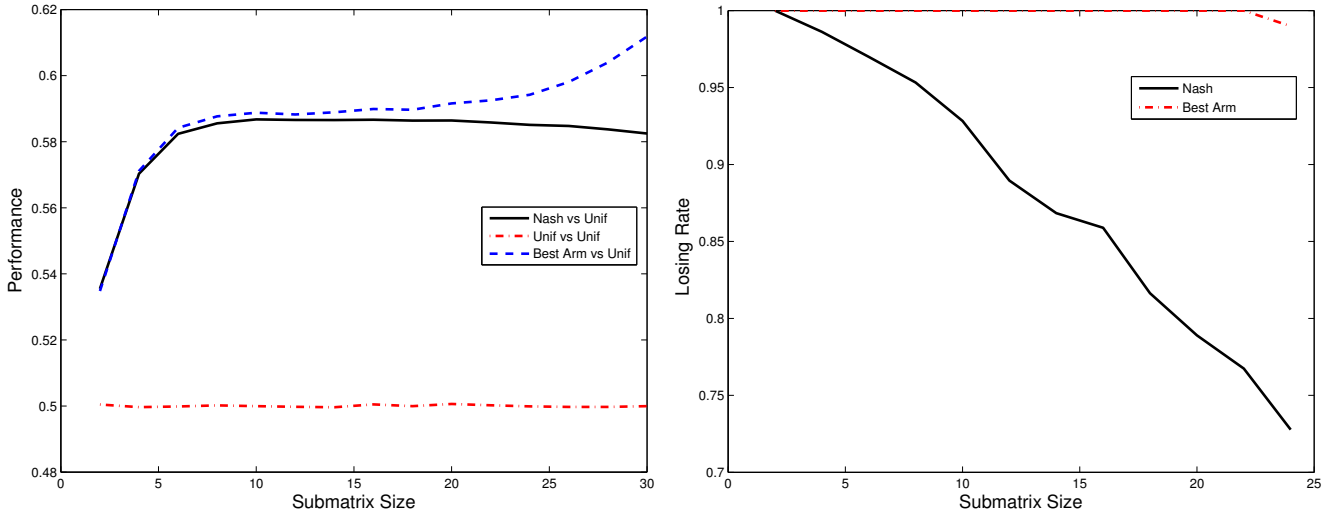

(a) Winning rate of 2 offline portfolios (namely Nash and $Best\ Arm$) against the uniform baseline, tested in generalization. X-axis: number K of policies considered in each portfolio. Y-axis: win rates. Experiments reproduced 10 000 times, standard deviations $< 10^{-2}$. Interpretation: we outperform (in generalization) the original GnuGo just by changing the probability distribution of random seeds.

(b) Losing rate of the $Nash$ and $Best\ Arm$ policies against the exploiter ($M = 32 - K$). X-axis: number K of considered seeds in the learning phase. Y-axis: loss rates. Experiments reproduced 100 times. We see that a simple learning (the "exploiter" easily crushes $Best\ Arm$, whereas $Nash$ resists to this difficult setting, in particular with a large learning set (i.e. a large learning matrix - and the rightmost point corresponds to 24 seeds, which therefore requires 576 games for training).

Fig. 2: **Game of Go: performance against the original GPP (left) and exploitability of Nash and BestArm respectively**.

reproduced 10 000 times (5 000 times for the Black player and 5 000 times for the White player) and standard deviations are smaller than 10^{-2}. Figure 3(b) shows the difference between a $Nash$ approach and a $Best\ Arm$ approach in terms of exploitability.

From Figure 3(a) we can observe, as it was the case in the game of Go, that there is a clear advantage to use either the $Nash$ or the $Best\ Arm$ approach when facing a new set of policies. As the size of the initial matrix grows, the winning rates of both $Nash$ and $Best\ Arm$ increase, in generalization, when compared to the baseline. Also, we observe that the shape of the curve for the $Nash$ approach is quite similar to the one seen in the game of Go. However, the $Best\ Arm$ approach keeps increasing almost linearly throughout the entire x-axis.

From Figure 3(b) it clearly appears that $Best\ Arm$ is a strategy very easy to exploit. Thus, while Figure 3(a) shows that the use of the $Best\ Arm$ approach outperforms $Nash$ versus the uniform baseline, $Nash$ is a much more resilient strategy.

Havannah: Figure 4 summarizes the results for the game of Havannah. Figure 4(a) presents the results of 2 offline portfolio algorithms (namely Nash and $Best\ Arm$) versus the uniform baseline. Same setting as for chess (number of experiments and same bound on the standard deviation).Figure 4(b) shows the difference between a $Nash$ approach and a $Best\ Arm$ approach in terms of exploitability.

From Figure 4(a) we can observe, as it was the case in the game of Go, that there is a clear advantage to use either the $Nash$ or the $Best\ Arm$ approach when facing a new set of policies. As the size of the initial matrix grows, the winning rates of both $Nash$ and $Best\ Arm$ increase when compared to the baseline.

From Figure 4(b) it clearly appears that $Best\ Arm$ is a strategy very easy to exploit. Thus, even if Figure 4(a) shows that the use of the $Best\ Arm$ approach outperforms $Nash$ versus the uniform baseline, $Nash$ is a much more resilient strategy.

The performance of $Nash$ and more especially $Best\ Arm$ increase significantly as the size of the submatrix grows. This is in sharp contrast with the 2 previous games. In the case of Havannah, the sharpest gain is towards the end of the x-axis, which suggests that further gains would be possible with bigger matrix.

Batoo: Figure 5 summarizes the results for the game of Batoo. Figure 5(a) presents the results of 2 different approaches versus the uniform baseline. Same setting as for Chess and Havannah.Figure 5(b) shows the difference between a $Nash$ approach and a $Best\ Arm$ approach in terms of exploitability. The x-axis represents the number of policies considered. The y-axis shows the loss rates. All experiments are reproduced 100 times.

From Figure 5(a) we can observe, as it was the case in the game of Go, that there is a clear advantage to use either the $Nash$ or the $Best\ Arm$ approach when facing a new set of policies. As the size of the initial matrix grows, the winning rates (in generalization) of both $Nash$ and $Best\ Arm$ increase when compared to the baseline.

From Figure 5(b) it clearly appears that $Best\ Arm$ is a strategy very easy to exploit. Thus, though Figure 5(a) shows that the use of the $Best\ Arm$ approach outperforms $Nash$ versus the uniform baseline, $Nash$ is a much more resilient strategy.

Conclusion: The performance of $Nash$ and $Best\ Arm$

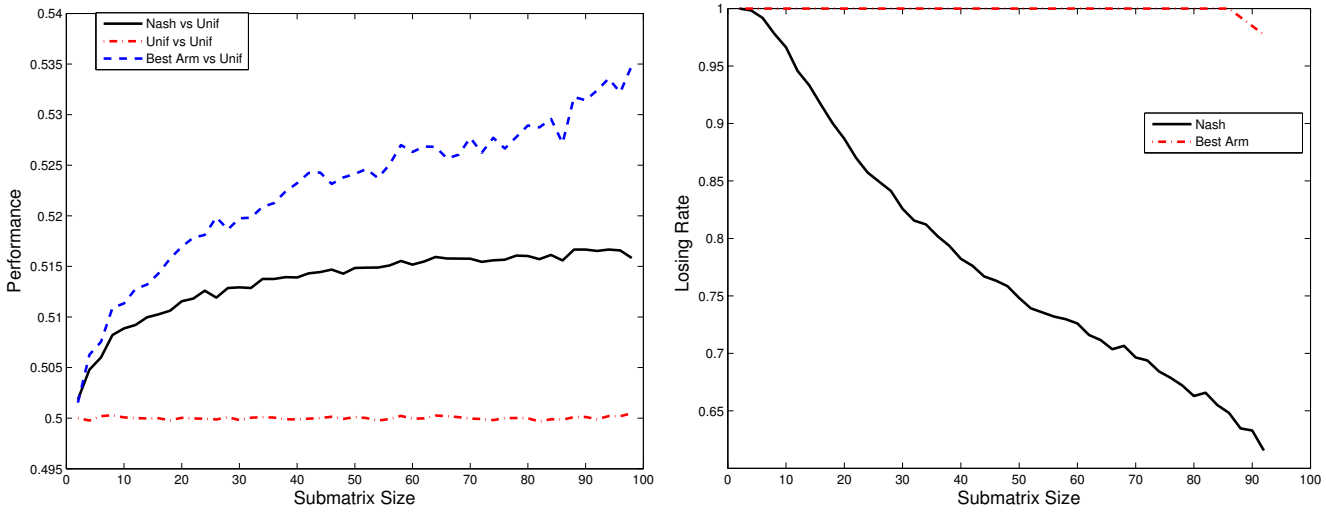

(a) Winning rate of 2 offline portfolios (Nash and $Best\ Arm$) against the uniform baseline in terms of generalization ability. Axes, number of experiments and standard deviation as in Fig. 2. We see that we have, for this Chess playing program, obtained a portfolio which is better than the original algorithm, just by modifying the distribution of random seeds.

(b) Losing rate of the $Nash$ and $Best\ Arm$ policies against the exploiter ($M = 100 - K$). Same axes as Fig. 2.

Fig. 3: **Game of Chess: performance against the original GPP (left) and exploitability of Nash and BestArm respectively**.

increase steadily as the size K of the submatrix grows. Also, we observe a behavior similar to the game of Go. The simultaneous action nature of the first move does not seem to impact the general efficiency of our approach.

B. Learning Online

The purpose of this section is twofold:

- Propose an adaptive algorithm, built automatically buy the random seed trick as in the case of $Nash$-Portfolio.
- Show the resilience of our offline-learning algorithms, namely $Nash$-Portfolio and $Best\ Arm$, against this adaptive algorithm - in particular, this shows a weakness of $Best\ Arm$ in terms of exploitability/overfitting.

Here we present the losing rate of UCBT (see Section III-B) against 3 baselines. The purpose is to evaluate whether learning a strategy online against a specific unknown opponent (baselines) can be efficiently done.

The first baseline is the Nash equilibrium (label $Nash$ and previously defined in Section III. The second baseline is the uniform player (label $Unif$) which consists in playing uniformly each option of the bandit. The third baseline consists in playing a single deterministic strategy (only one random seed) regardless of the opponent.

Go: Figure 6(a) (and Figure 6(b)) shows the learning of UCBT for the Black player (and White respectively) for the game of Go.

First and foremost, as the number of iterations grows, there is a clear learning against both $Nash$ and $Unif$ baselines. We see that (i) UCBT eventually reaches, against Nash-portfolio, approximately the value of the game for each player, (ii) the Nash-portfolio is among the most difficult opponents (the curve decreases slowly only). We can also observe from Figures 6(a) and 6(b) that against the $Unif$ baseline UCBT learns a strategy that outperforms this opponent.

When it plays as the Black player, it takes less than 2^7 (128) games to learn the correct strategy and win with a 100 % ratio against every single deterministic variant. As the White player, it is even faster with only 2^5 games required to always win. Also, it is without surprise that the losing rate is lower when UCBT is the first player.

Chess: Figure 6(c) (and Figure 6(d)) shows the learning of UCBT for the Black player (and White respectively) for the game of Chess.

Again, as the number of iterations grows, there is a clear learning against both $Nash$ and $Unif$ baselines. UCBT eventually reaches, against Nash-portfolio, almost the value of the game for each player. Moreover, by looking at the slope of the curves, we see that the Nash-portfolio is among the most difficult opponents. We can also observe from Figures 6(c) and 6(d) that against the $Unif$ baseline UCBT learns a strategy that outperforms this opponent. This is consistent with the theory behind UCBT.

When it plays as the Black player, it takes less than 2^7 games to learn the correct strategy and win with a 100 % ratio against every single deterministic variant. As the White player, it is even faster with only 2^6 games required to always win. In Section V-A we observe that the uniform strategy for the game of Chess is much more difficult to play against than the uniform strategy for the game of Go. Figures 6(c) and 6(c) corroborate this results as the slope of learning against the uniform strategy is less pronounced in Chess than in Go.

Havannah: Figure 7(a) (resp. Figure 7(b)) shows the learning of UCBT for the Black player (resp. White player) for the game of Havannah.

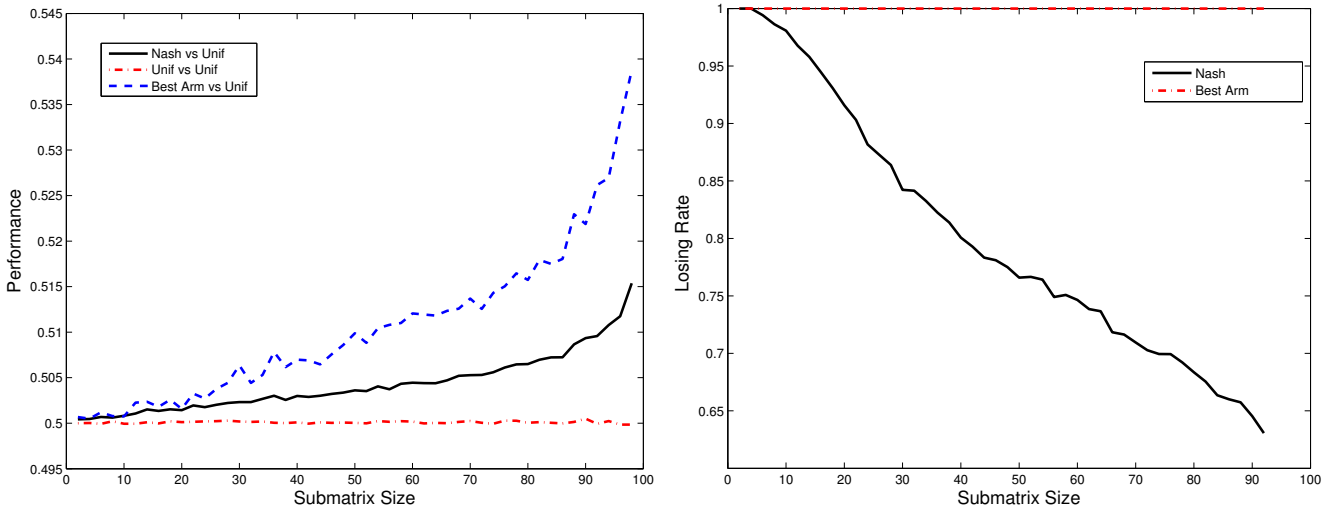

(a) Winning rate of 2 offline portfolios (Nash and $Best\ Arm$) against the uniform baseline in generalization. Same setting as in Fig. 2. We see that we get a program which outperforms the original Havannah artificial intelligence just by changing the probability distribution of random seeds.

(b) Losing rate of the $Nash$ and $Best\ Arm$ policies against the exploiter ($M = 100 - K$). Same setting as in Fig. 2.

Fig. 4: **Game of Havannah: performance against the original GPP (left) and exploitability of Nash and BestArm respectively**.

Once more, as the number of iterations grows, there is a clear learning against both $Nash$ and $Unif$ baselines. Moreover, by looking at the slope of the curves, we see that the Nash-portfolio is harder to exploit than other opponents, and in particular than the original algorithm, i.e. the uniform random seed. We can also observe from Figures 7(a) and 7(b) that against the $Unif$ baseline UCBT learns a strategy that outperforms this opponent. However, it takes about 2^6 iterations before the learning really kicks in.

When it plays as the Black player, it takes less than 2^5 games to learn the correct strategy and win with a 100 % ratio against every single deterministic variant. As the White player, it is even faster with only 2^5 games required to always win.

Batoo: Figure 7(c) and Figure 7(d) show the learning of UCBT for the Black and White players respectively for the game of simplified Batoo.

Even though this game contains a critical simultaneous action at the beginning, the results are quite similar to the previous games. As the number of iterations grows, there is a clear learning against both $Nash$ and $Unif$ baselines. Moreover, by looking at the slope of the curves, we see that the Nash-portfolio is among the most difficult opponents - it is harder to exploit than the original algorithm with uniform seed. We can also observe from Figure 7 that against the $Unif$ baseline UCBT learns a strategy that outperforms this opponent.

When it plays as the Black player, it takes less than 2^7 games to learn the correct strategy and win with a 100 % ratio against every single deterministic variant. As the White player, it is even faster with only 2^7 games required to always win.

We now switch to UCBT applied to the Variants problem. The losing rates of the recommended variant are presented in Fig. 8.

First and foremost, as the number of iterations grows, there is a clear learning against both $Nash$ and $Unif$ baselines. We see that (i) UCBT eventually reaches, against Nash-portfolio, approximately the value of the game for each player (ii) the Nash-portfolio is among the most difficult opponents (the curve decreases slowly only). We can also observe from Figure 8 that against the $Unif$ baseline UCBT learns a strategy that outperforms his opponent.

1) Conclusions: UCBT can learn very efficiently against a fixed deterministic opponent; this confirms its ability for eTeaching - a human opponent can learned her weaknesses by playing against a UCBT program. UCBT, after learning, performs better than Nash-portfolio against Uniform, showing that even against a stochastic opponent it can perform well, and in particular better than the Nash. This is not a contradiction with the Nash optimality; the $Nash$ portfolio is optimal in an agnostic sense, whereas UCBT tries to overfit its opponent and can therefore exploit it better.

2) Generalization ability of online portfolios: We validated offline portfolios both against the GPPs used in the training, and against other GPPs. For online learning, the generalization ability does not have the same meaning, because online learning is precisely aimed at exploiting a given opponent. Nonetheless, we can consider what happens if we online learn random seeds against the uniform portfolio, and then play games against the original GPP.

The answer can be derived mathematically. From the consistency of UCBT, we deduce that UCBT-portfolio, against a randomized seed, will converge to $Best\ Arm$. Therefore, the asymptotic winning rate of UCBT-portfolio when learning

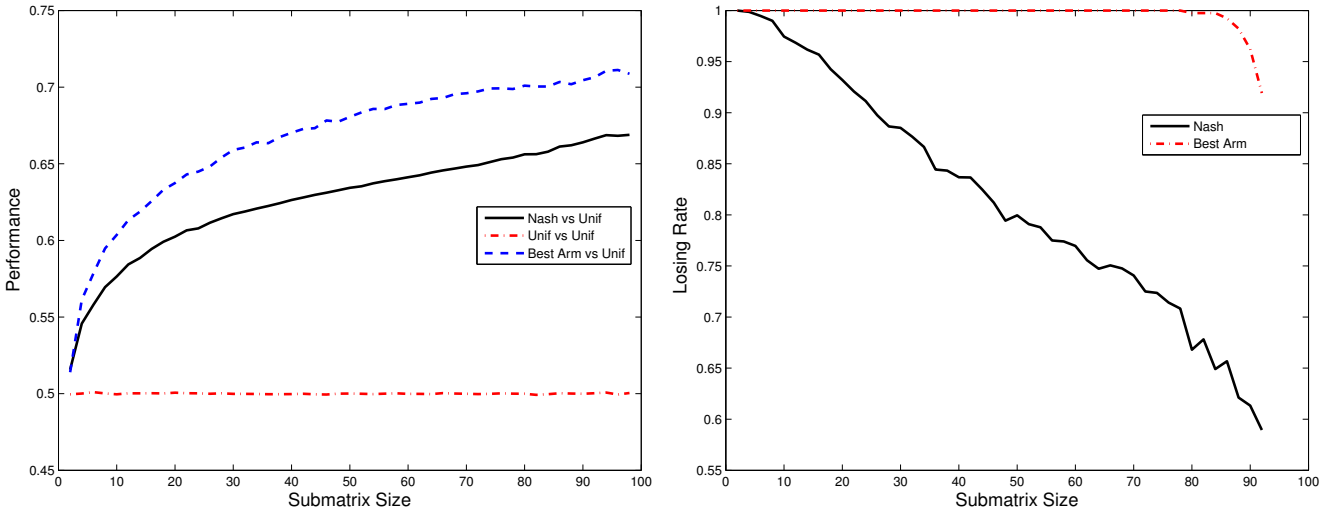

(a) Winning rate of 2 offline portfolios against the uniform baseline in generalization. We see that we have obtained a version of our Batoo playing program which outperforms the original program, just by modifying the probability distribution over random seeds.

(b) Losing rate of the $Nash$ and $Best\ Arm$ policies against the exploiter ($M = 100 - K$). Same setting as in Fig. 2.

Fig. 5: **Game of Batoo: performance against the original GPP (left) and exploitability of Nash and BestArm respectively**.

against the original GPP, using a training against a fixed number of random seeds, is the same as shown for $Best\ Arm$ in Section V-A: 62% in Go, 54% in Havannah, 53.5% in Chess, 71% in Batoo. In the case of Batoo we see that this generalization success rate is better than the empirical success rate from Fig. 5; this is not surprising as we consider the asymptotic success rate whereas we clearly see on Figure 5 that the asymptotic rate is not yet reached.

VI. Robustness: the transfer to other opponents

Results above were performed in a classical machine learning setting, i.e. with cross-validation; we now check the transfer, i.e. the fact that we improve a GPP not only in terms of winning rate against the baseline version, but also in terms of better performance when we test its performance

- by playing against another, distinct, GPP;
- by analysis with a reference GPP, stronger thanks to huge thinking time.

This means, that whereas previous sections have obtained results such as

"When our algorithm takes A as baseline GPP, the boosted counterpart A' outperforms A by XXX % winning rate. (with XXX>50%)"

we get results such as:

"When our algorithm takes A as baseline GPP, the boosted counterpart A' outperforms A in the sense that the winning rate of A' against B is greater than the winning rate of A against B, for each B in a family { B1, B2, ..., Bk } of programs different from A."

A. Transfer to GnuGo

We applied BestArm to GnuGo, a well known AI for the game of Go, with Monte Carlo tree search and a budget of 400 simulations. The BestArm approach was applied with a 100x100 learning matrix, corresponding to seeds $\{1, \ldots, 100\}$ for Black and seeds $\{1, \ldots, 100\}$ for White.

TABLE II: Performance (winning rate) of BestArm-Gnugo-MCTS against various GnuGo-default programs, compared to the performance of the default Gnugo-MCTS. The results are for GnuGo-MCTS playing as Black vs GnuGo-classical playing as White, and the games are completely independent of the learning phase - which use only Gnugo-MCTS. Results are averaged over 1000 games. All results in 5x5, komi 6.5, with a learning over a 100x100 matrix of games played between 100 random seeds for Black and 100 random seeds for White.

Opponent	Performance of BestArm	Performance of the original algorithm with randomized random seed
GnuGo-classical level 1	1. (± 0)	.995 (± 0.002)
GnuGo-classical level 2	1. (± 0)	.995 (± 0.002)
GnuGo-classical level 3	1. (± 0)	.99 (± 0.002)
GnuGo-classical level 4	1. (± 0)	1. (± 0)
GnuGo-classical level 5	1. (± 0)	1. (± 0)
GnuGo-classical level 6	1. (± 0)	1. (± 0)
GnuGo-classical level 7	.73 (± .013)	.061 (± .004)
GnuGo-classical level 8	.73 (± .013)	.106 (± .006)
GnuGo-classical level 9	.73 (± .013)	.095 (± .006)
GnuGo-classical level 10	.73 (± .013)	.07 (± .004)

Then, we tested the performance against GnuGo "classical", i.e. the non-MCTS version of GnuGo; this is a really different AI with different playing style. We got positive results as shown in Table II. Results are presented for Black; for White the BestArm had a negligible impact.

B. Transfer: validation by a MCTS with long thinking time

Figure 9 provides a summary of differences between moves chosen (at least with some probability) by the original algorithm, and the ones chosen in the same situation by the algorithm with optimized seed. These situations are the 8 first

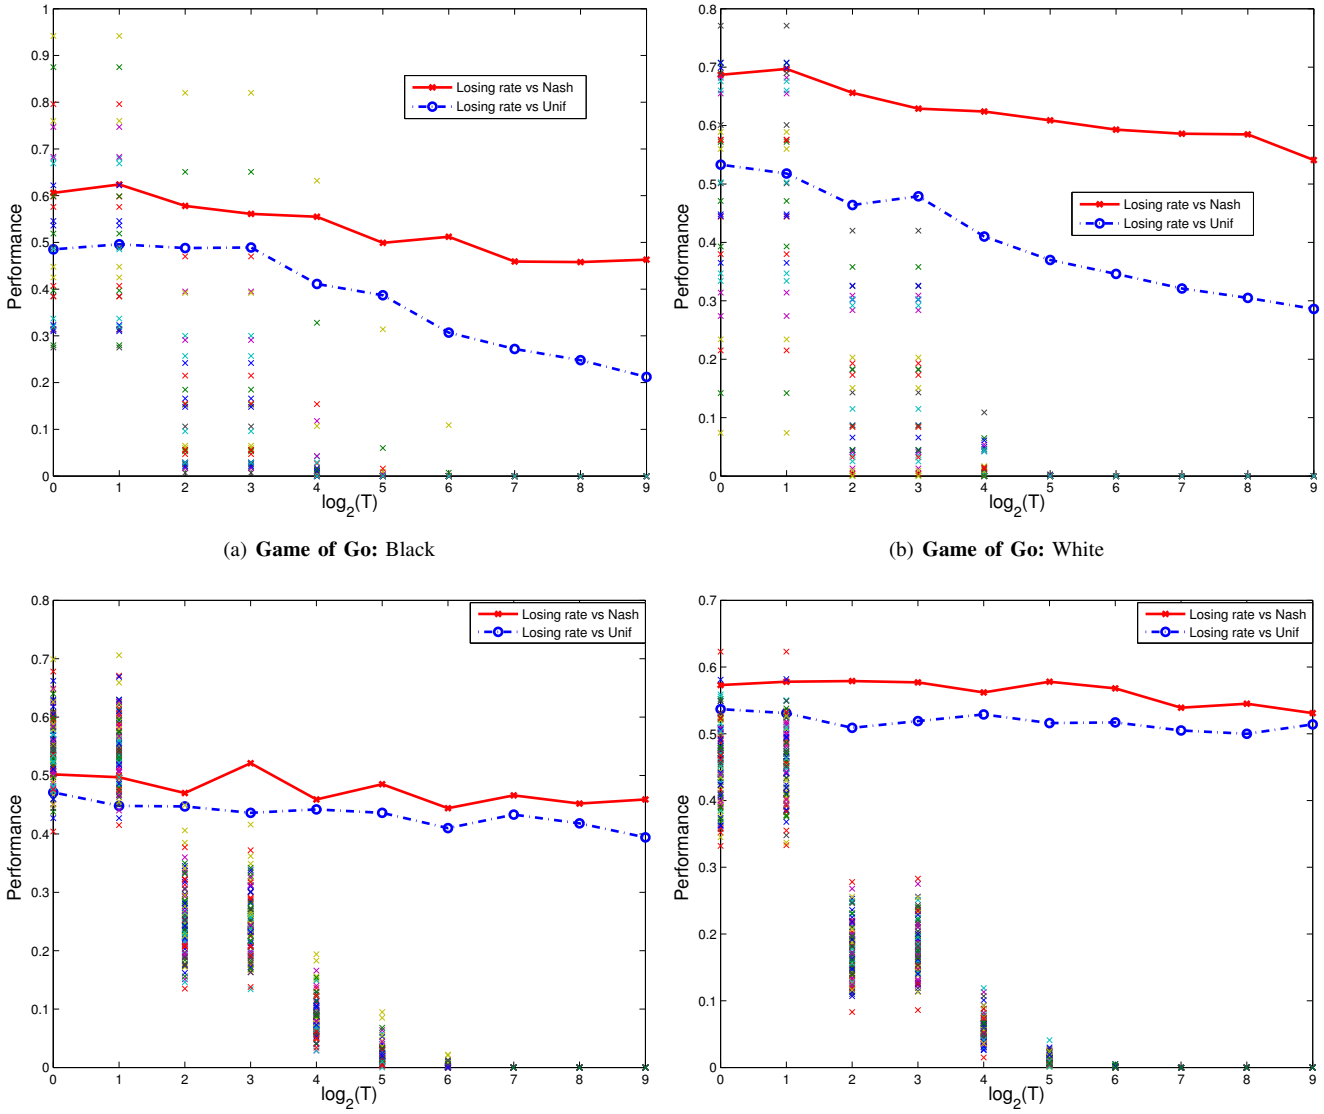

(a) **Game of Go:** Black

(b) **Game of Go:** White

(c) **Game of Chess:** White

(d) **Game of Chess:** Black

Fig. 6: **Game of Go and Chess**. Losing rate of UCBT-portfolio, versus the online learning time, for (i) Nash-Portfolio (red line) (ii) Uniform portfolio (dotted blue line) (iii) each option independently (stars). X-axis: log2(number of iterations of UCBT (i.e. number of played games for learning). Y-axis: frequency at which the game is lost. Experiments reproduced 1 000 times. Standard deviations $\leq 10^{-4}$. Learning is visible in the sense that curves essentially decrease.

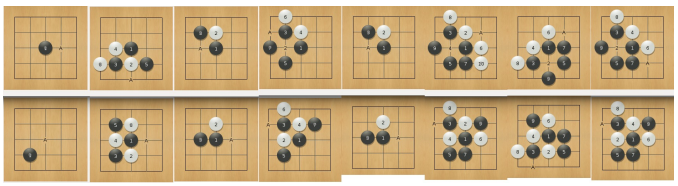

Fig. 9: Comparison between moves played by BestArm-MCTS (top) and the original MCTS algorithm (bottom) in the same situations.

differences between games played by the original GnuGo and by the GnuGo with our best seed.

We use GnugoStrong, i.e. Gnugo with a larger number of simulations, for checking if Seed 59 leads to better moves.

GnugoStrong is precisely defined as "gnugo –monte-carlo –mc-games-per-level 100000 –level 1". We provide below some situations in which Seed 59 (top) proposes a move different from the original Gnugo with the same number of simulations. Gnugo is not deterministic; therefore this is simple the 8 first differences found in our sample of games (we played games until we find 8 differences).

We consider that GnugoStrong concludes that a situation is a win (resp. loss) if, over 5 games played from this situation, we always get a win (resp. loss). The conclusions from this GnugoStrong experiment (8 situations) are as follows, for the 8 situations above respectively:

1) GnugoStrong prefers Top; Bottom is considered as a

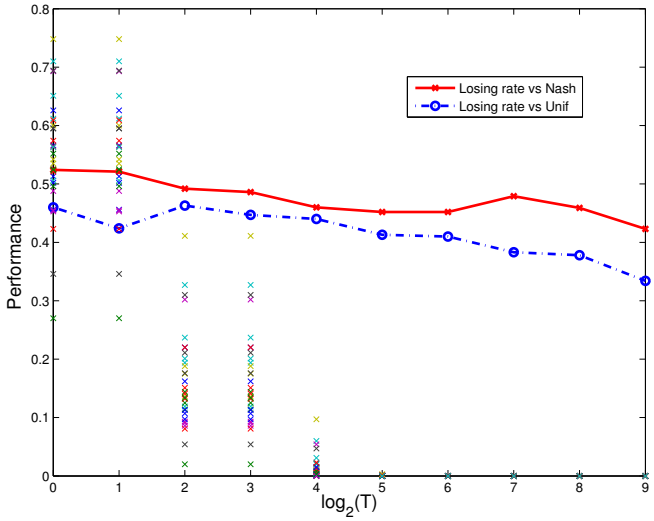

(a) **Game of Havannah:** Black

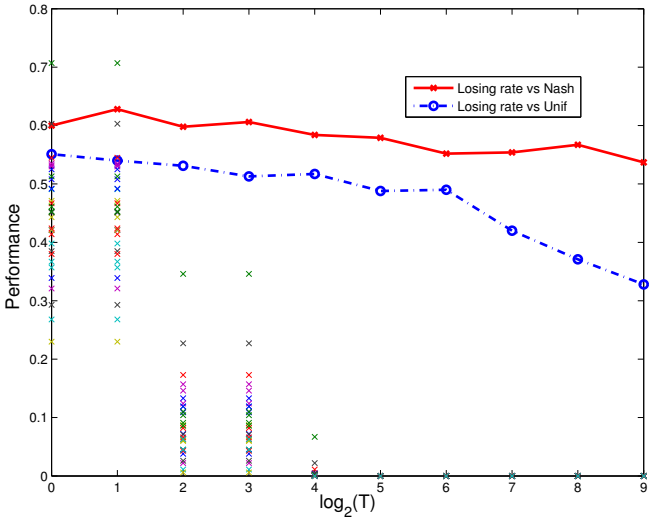

(b) **Game of Havannah:** White

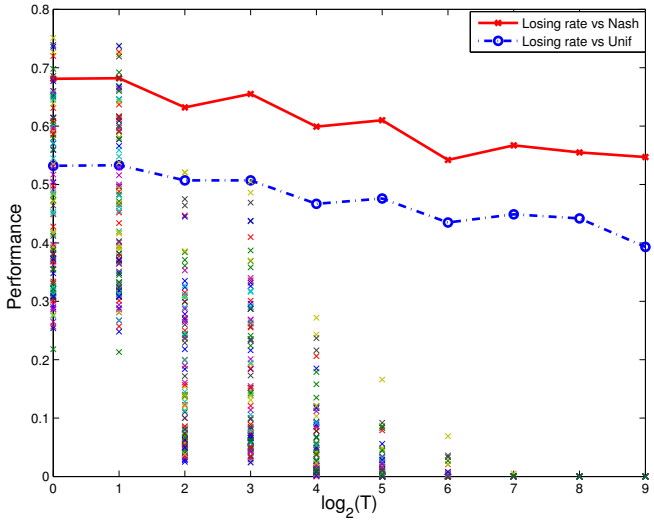

(c) **Game of Batoo:** Black

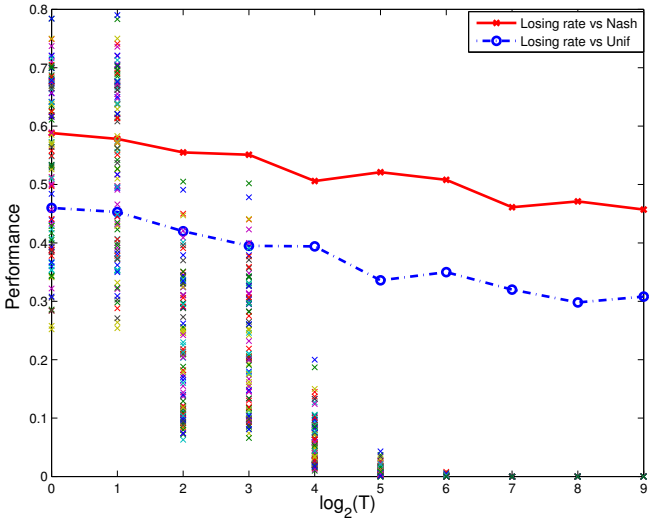

(d) **Game of Batoo:** White

Fig. 7: **Game of Havannah and Batoo**. Losing rate of UCBT-portfolio, versus the online learning time, against (i) Nash-Portfolio (red line) (ii) Uniform portfolio (dotted blue line) (iii) each option independently (stars). X-axis: log2(number of iterations of UCBT (i.e. number of played games for learning). Y-axis: frequency at which the game is lost. Experiments reproduced 1 000 times, standard deviations $\leq 10^{-4}$. Learning is visible in the sense that curves essentially decrease.

loss.

2) Here black moves are the same up to symmetries. Both are considered as wins for Black.
3) Both choices are considered as wins for Black.
4) Both choices are considered as wins for Black.
5) Both choices are considered as wins for Black.
6) GnugoStrong prefers Top; Bottom is considered as a loss.
7) GnugoStrong prefers Top; Bottom is considered as a loss.
8) GnugoStrong prefers Top; Bottom is considered as a loss.

As a conclusion, in 4 cases, GnugoStrong prefers the move chosen by the modified MCTS (with seed chosen by BestArm). In 4 cases, moves are equivalent.

VII. Additional experiments: simpler tools

We have seen that the Nash portfolio has some advantages compared to the BestArm method, namely robustness against a learning opponent. On the other hand, the Nash mechanism induces a computational overhead and some implementation complexity. Therefore, we here investigate the relative performance of different methods:

- Nash portfolio.
- BestArm portfolio.
- BestHalf, which is a uniform random choice among the options with performance better than the median performance over the learning set. In other words, given

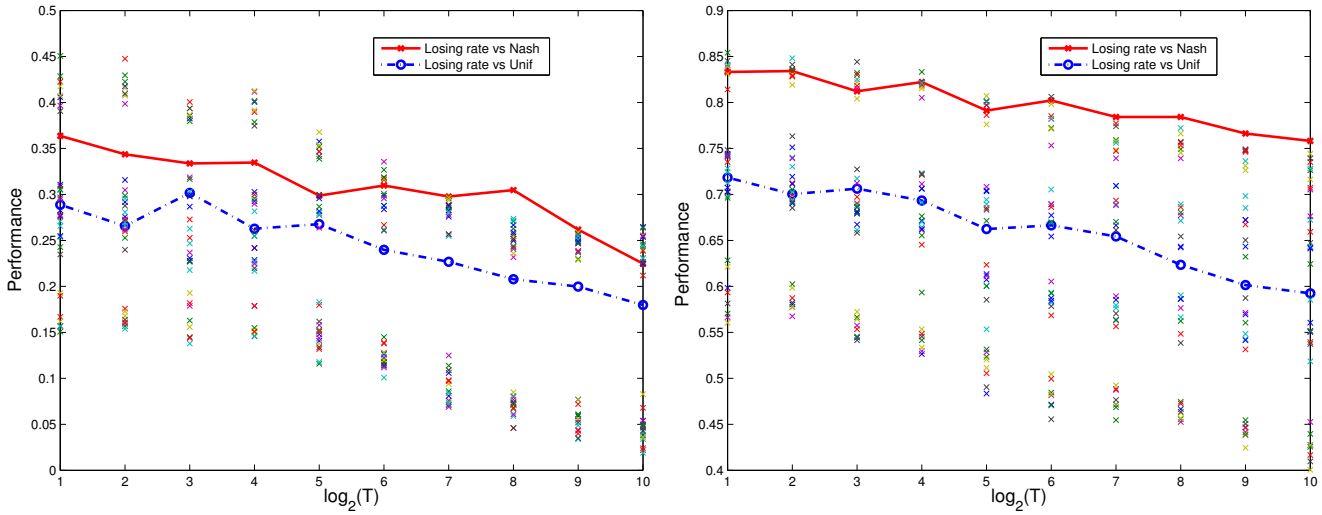

(a) Black (b) White

Fig. 8: **Game of Go using different variants**. Losing rate of UCBT-portfolio, versus the online learning time, against (i) Nash-Portfolio (red line) (ii) Uniform portfolio (dotted blue line) (iii) each option independently (stars). X-axis: log2(number of iterations of UCBT (i.e. number of played games for learning). Y-axis: frequency at which the game is lost. Experiments reproduced 1 000 times, standard deviations $\leq 10^{-4}$. Learning is visible in the sense that curves essentially decrease. We see that deterministic variants are quickly crushed. We see that the $Nash$ portfolio resists better than the uniform portfolio.

the matrix M, it selects the rows i with sum $\sum_{j=1}^{K} M_{i,j}$ above the median. These options (up to ties and rounding, there are $K/2$ such options) are played with the same probability.

- Uniform, i.e. randomly choosing the option.
- Exploiter, who "cheats" by choosing the best element in its portfolio *while knowing its performance against its opponent*. By definition, Exploiter can not play against itself.

The four first are policies which can be used in real life. The fifth one requires offline training, which is difficult in a competition unless your opponent has shared his program.

All results are obtained with learning matrices distinct for Black and White so that there is no overfitting bias.

Results are provided in Fig. 10 for Chess (portfolio of seeds) and Go (portfolio of variants), 11 for Havannah and Batoo (portfolio of seeds). As a conclusion:

- BestHalf is a good solution, performing similarly to Nash, including in terms of results against Exploiter, whereas it is very simple.
- The case of "variants" has non binary outputs; and, consequently, it is less prone to overfitting; even BestArm performs reasonably well against Exploiter.

VIII. Conclusion

We proposed three algorithms for combining policies:

- Three offline algorithms, namely the Nash-Portfolio, the simpler BestHalf-portfolio which approximates Nash-Portfolio quite well, and the $Best\ Arm$ Portfolio. Given a family of algorithms, or a single algorithm "exploded" into several algorithms (by the use of random seeds or by varying its parameterizations), they build, offline, a new algorithm.
- The Bandit(UCBT)-Portfolio, which learns online, given an opponent; this one is adaptive.

We tested them on Go, Chess, Havannah, Batoo. Other publications include experiments on Domineering, Atari, Breakthrough [41]. A work on Tsumego has investigated the use of seeds learnt online (see results comparing MCTS(1) to other methods in [42]).

We have seen that:

- The Nash-Portfolio is more diversified than any of its components in the sense that it is harder to learn against (i.e. harder to exploit) than any of its components and harder to learn against than the uniform-Portfolio. $Best\ Arm$ is less resilient (easy to exploit, converging to 0% success rate against UCBT portfolio), though quite efficient against the uniform-Portfolio.
- The UCBT-Portfolio can learn a combination until reaching optimal exploitation of a stationary opponent (this is mathematically guaranteed by properties of UCBT, which is a consistent bandit algorithm in the discrete setting). In particular, it defeated clearly each deterministic variants in Fig. 6, reaching 100% winning rate. It also performs quite well against the default policy, which is uniformly randomized random seed. This shows that we can enhance a randomized algorithm a lot, just by biasing the choice of the random seed. The Nash-Portfolio resists much better (than uniform or $Best\ Arm$) to a UCBT-portfolio, showing that the $Nash$ portfolio is hard to exploit. $Best\ Arm$ is, by construction, the asymptotic limit of UCBT learning against the uniform portfolio, and against the original algorithm.

The Nash-portfolio and $Best\ Arm$ portfolio perform well against the original GPP (Section V-A). Therefore our tools provide an easy improvement on top of randomized algorithms, or on top of variants of an algorithm. Results in Section VI show the robustness of the results, in the sense that the improved GPP is not only winning against the original GPP, but it also has a better winning rate against other opponents than the ones used in the learning.

The computational cost could potentially become an issue when combining parameterizations (in Section IV-B - with large computational cost because we averaged multiple games for building the matrix). It is not the case for combining random seeds (Section IV-A), where deterministic games are used and therefore there is no point in duplicating games. It should also be pointed out that solving Nash is fast, and if needs be, we can further the speed of the computation using algorithms that can ϵ-approximate a NE in sublinear time [30], [31] - so that a number of games linear in $\max(K, K')\log(KK')/\epsilon^2$ (i.e. far less than the number $K \times K'$ of elements in the matrix) is sufficient for a precision ϵ. Importantly, even in the case of combining variants, all computational costs are offline, so that there is no significant online computational overhead.

Results are threefold:

- An improvement in terms of playing strength measured in direct games, as our $Best\ Arm$, the $Best\ Half$ policy suggested by a reviewer, and our $Nash$-Portfolio all outperform the original randomized algorithm.
- An improvement in terms of eTeaching (use of our program as a teaching tool); our UCBT-Portfolio algorithm is adaptive and difficult to overfit. This does not involve any additional computational power as UCBT is online and has a negligible internal cost. This makes our UCBT tool suitable for online learning.
- An improvement in terms of resilience; for learning size at least 40 seeds, our Nash-Portfolio is harder to overfit than any of its components and than the uniform portfolio (in Fig. 6 - 8). At the end of the offline computation of the Nash equilibrium, it is just a bias in the random seed distribution, so the additional computational cost is negligible. The non-intuitive key point in the "random seed" part of this work is that biasing the random seed has an impact.

Further work.

The easiest and most promising further work consists in using fast approximate Nash equilibria, so that the full matrix does not have to be computed. This should extend by far the number K of arms that our method can handle. Another further work is the use of discounted bandits for the UCBT bandit, so that very old games have little influence on the current games - this should bring improvements in terms of adaptivity, when playing against a non-stationary opponent such as a human. An infinite set of seeds should be considered in UCBT, so that exact optimality might, asymptotically, be reached. An adapted bandit algorithm already exists for such cases [43]. We have considered random seeds for whole runs; more specialized representations are under work. The present work might indeed be used for deterministic algorithms without any parameterization: we might e.g. randomize the parameters of an evaluation function used in an alpha-beta algorithm.

Acknowledgements

We are grateful to the Dagstuhl seminar for interesting discussions around coevolution.

REFERENCES

[1] P. E. Utgoff, "Perceptron trees: A case study in hybrid concept representations," in *National Conference on Artificial Intelligence*, 1988, pp. 601–606.

[2] D. W. Aha, "Generalizing from case studies: A case study," in *Proceedings of the 9th International Workshop on Machine Learning*. Morgan Kaufmann Publishers Inc., 1992, pp. 1–10.

[3] E. Nudelman, K. Leyton-Brown, H. H. Hoos, A. Devkar, and Y. Shoham, "Understanding random sat: beyond the clauses-to-variables ratio," in *Principles and Practice of Constraint Programming CP 2004*, M. Wallace, Ed., vol. 3258 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2004, pp. 438–452.

[4] L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Hydra-mip: automated algorithm configuration and selection for mixed integer programming," in *RCRA Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion at the International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.

[5] L. Kotthoff, "Algorithm selection for combinatorial search problems: A survey," *AI Magazine*, vol. 35, no. 3, pp. 48–60, 2014. [Online]. Available: http://www.aaai.org/ojs/index.php/aimagazine/article/view/2460

[6] ——, "Algorithm selection for combinatorial search problems: A survey," *AI Magazine*, vol. 35, no. 3, pp. 48–60, 2014.

[7] A. R. KhudaBukhsh, L. Xu, H. H. Hoos, and K. Leyton-Brown, "Satenstein: Automatically building local search SAT solvers from components," in *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, 2009, pp. 517–524. [Online]. Available: http://ijcai.org/papers09/Papers/IJCAI09-093.pdf

[8] F. Hutter, H. H. Hoos, K. Leyton-Brown, and T. Stützle, "Paramils: An automatic algorithm configuration framework," *J. Artif. Int. Res.*, vol. 36, no. 1, pp. 267–306, Sep. 2009. [Online]. Available: http://dl.acm.org/citation.cfm?id=1734953.1734959

[9] D. Bolme, J. Beveridge, B. Draper, P. Phillips, and Y. Lui, "Automatically searching for optimal parameter settings using a genetic algorithm," in *Computer Vision Systems*, ser. Lecture Notes in Computer Science, J. Crowley, B. Draper, and M. Thonnat, Eds. Springer Berlin Heidelberg, 2011, vol. 6962, pp. 213–222. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-23968-7_22

[10] B. Bouzy, M. Métivier, and D. Pellier, "Hedging algorithms and repeated matrix games," in *ECML Workshop on Machine Learning and Data Mining In and Around Games*, 2011.

[11] M. Swiechowski and J. Mandziuk, "Self-adaptation of playing strategies in general game playing," *IEEE Trans. Comput. Intellig. and AI in Games*, vol. 6, no. 4, pp. 367–381, 2014. [Online]. Available: http://dx.doi.org/10.1109/TCIAIG.2013.2275163

[12] J. Borrett and E. Tsang, "Towards a formal framework for comparing constraint satisfaction problem formulations," University of Essex, Colchester, UK, Tech. Rep. CSM-264, 1996.

[13] V. Vassilevska, R. Williams, and S. L. M. Woo, "Confronting hardness using a hybrid approach," in *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. ACM, 2006, pp. 1–10.

[14] L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Satzilla: Portfolio-based algorithm selection for sat." *J. Artif. Intell. Res.(JAIR)*, vol. 32, pp. 565–606, 2008.

[15] P. O. Stalph, M. Ebner, M. Michel, B. Pfaff, and R. Benz, "Multiobjective evolution of a fuzzy controller in a sewage treatment plant," in *GECCO*, C. Ryan and M. Keijzer, Eds., pp. 535–536.

[16] S. Teraoka, T. Ushio, and T. Kanazawa, "Voronoi coverage control with time-driven communication for mobile sensing networks with obstacles." in *CDC-ECE*. IEEE, 2011, pp. 1980–1985. [Online]. Available: http://dblp.uni-trier.de/db/conf/cdc/cdc2011.html#TeraokaUK11

[17] O. Lejri and M. Tagina, "Representation in case-based reasoning applied to control reconfiguration," in *Advances in Data Mining. Applications and Theoretical Aspects*, ser. Lecture Notes in Computer Science, P. Perner, Ed. Springer Berlin Heidelberg, 2012, vol. 7377, pp. 113–120. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-31488-9_10

[18] D. L. Saint-Pierre and O. Teytaud, "Nash and the Bandit Approach for Adversarial Portfolios," in *CIG 2014 - Computational Intelligence in Games*, ser. Computational Intelligence in Games, IEEE. Dortmund, Germany: IEEE, Aug. 2014, pp. 1–7. [Online]. Available: https://hal.inria.fr/hal-01077628

[19] R. Gaudel, J.-B. Hoock, J. Pérez, N. Sokolovska, and O. Teytaud, "A principled method for exploiting opening books," in *Computers and Games*. Springer, 2011, pp. 136–144.

[20] S. Kadioglu, Y. Malitsky, A. Sabharwal, H. Samulowitz, and M. Sellmann, "Algorithm selection and scheduling," in *17th International Conference on Principles and Practice of Constraint Programming*, 2011, pp. 454–469.

[21] M. Gagliolo and J. Schmidhuber, "Learning dynamic algorithm portfolios," vol. 47, no. 3-4, 2006, pp. 295–328.

[22] W. Armstrong, P. Christen, E. McCreath, and A. P. Rendell, "Dynamic algorithm selection using reinforcement learning," in *International Workshop on Integrating AI and Data Mining*, 2006, pp. 18–25.

[23] V. Nagarajan, L. S. Marcolino, and M. Tambe, "Every team deserves a second chance: Identifying when things go wrong (student abstract version)," in *29th Conference on Artificial Intelligence (AAAI 2015), Texas, USA*, 2015.

[24] R. A. Valenzano, N. R. Sturtevant, J. Schaeffer, K. Buro, and A. Kishimoto, "Simultaneously searching with multiple settings: An alternative to parameter tuning for suboptimal single-agent search algorithms," in *Proceedings of the 20th International Conference on Automated Planning and Scheduling, ICAPS 2010, Toronto, Ontario, Canada, May 12-16, 2010*, R. I. Brafman, H. Geffner, J. Hoffmann, and H. A. Kautz, Eds. AAAI, 2010, pp. 177–184. [Online]. Available: http://www.aaai.org/ocs/index.php/ICAPS/ICAPS10/paper/view/1457

[25] K. Hoki, T. Kaneko, A. Kishimoto, and T. Ito, "Parallel dovetailing and its application to depth-first proof-number search," *ICGA Journal*, vol. 36, no. 1, pp. 22–36, 2013.

[26] J. V. Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton University Press, 1944. [Online]. Available: http://jmvidal.cse.sc.edu/library/neumann44a.pdf

[27] J. Nash, "The bargaining problem," *Econometrica*, vol. 18, no. 2, pp. 155–162, April 1950.

[28] V. N. Vapnik, *The Nature of Statistical Learning*. Springer Verlag, 1995.

[29] K. Gale and Tucker, "Linear programming and the theory of games," in *Activity Analysis of Production and Allocation*, Koopmans, Ed. Wiley, 1951, ch. XII.

[30] M. D. Grigoriadis and L. G. Khachiyan, "A sublinear-time randomized approximation algorithm for matrix games," *Operations Research Letters*, vol. 18, no. 2, pp. 53–58, Sep 1995.

[31] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "Gambling in a rigged casino: the adversarial multi-armed bandit problem," in *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 1995, pp. 322–331.

[32] H. Samulowitz and R. Memisevic, "Learning to solve QBF," in *Proceedings of the 22nd National Conference on Artificial Intelligence*. AAAI, 2007, pp. 255–260.

[33] T. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, pp. 4–22, 1985.

[34] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2/3, pp. 235–256, 2002.

[35] J. Audibert, R. Munos, and C. Szepesvári, "Exploration-exploitation tradeoff using variance estimates in multi-armed bandits," *Theor. Comput. Sci.*, vol. 410, no. 19, pp. 1876–1902, 2009. [Online]. Available: http://dx.doi.org/10.1016/j.tcs.2009.01.016

[36] R. Coulom, "Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search," *In P. Ciancarini and H. J. van den Herik, editors, Proceedings of the 5th International Conference on Computers and Games, Turin, Italy*, pp. 72–83, 2006.

[37] B. Robertie, "Backgammon," *Inside Backgammon*, vol. 2, no. 1, p. 4, 1980.

[38] M. Campbell, A. J. Hoane Jr, and F.-h. Hsu, "Deep Blue," *Artificial intelligence*, vol. 134, no. 1, pp. 57–83, 2002.

[39] S. Gelly and D. Silver, "Combining online and offline knowledge in UCT," in *ICML '07: Proceedings of the 24th international conference on Machine learning*. New York, NY, USA: ACM Press, 2007, pp. 273–280.

[40] Tao Uct Sig. (2014) Tao games axis. [Online]. Available: http://www.lri.fr/~teytaud/games.html

[41] T. Cazenave, J. Liu, and O. Teytaud, "The rectangular seeds of domineering, atari-go and breakthrough," in *Computational Intelligence and Games (CIG), 2015 IEEE Congress on*. IEEE, 2015.

[42] D. L. St-Pierre, J. Liu, and O. Teytaud, "Nash reweighting of monte carlo simulations: Tsumego," in *Evolutionary Computation (CEC), 2015 IEEE Congress on*. IEEE, 2015, pp. 1458–1465.

[43] Y. Wang, J. Audibert, and R. Munos, "Algorithms for infinitely many-armed bandits," in *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, 2008, pp. 1729–1736. [Online]. Available: http://papers.nips.cc/paper/3452-algorithms-for-infinitely-many-armed-bandits

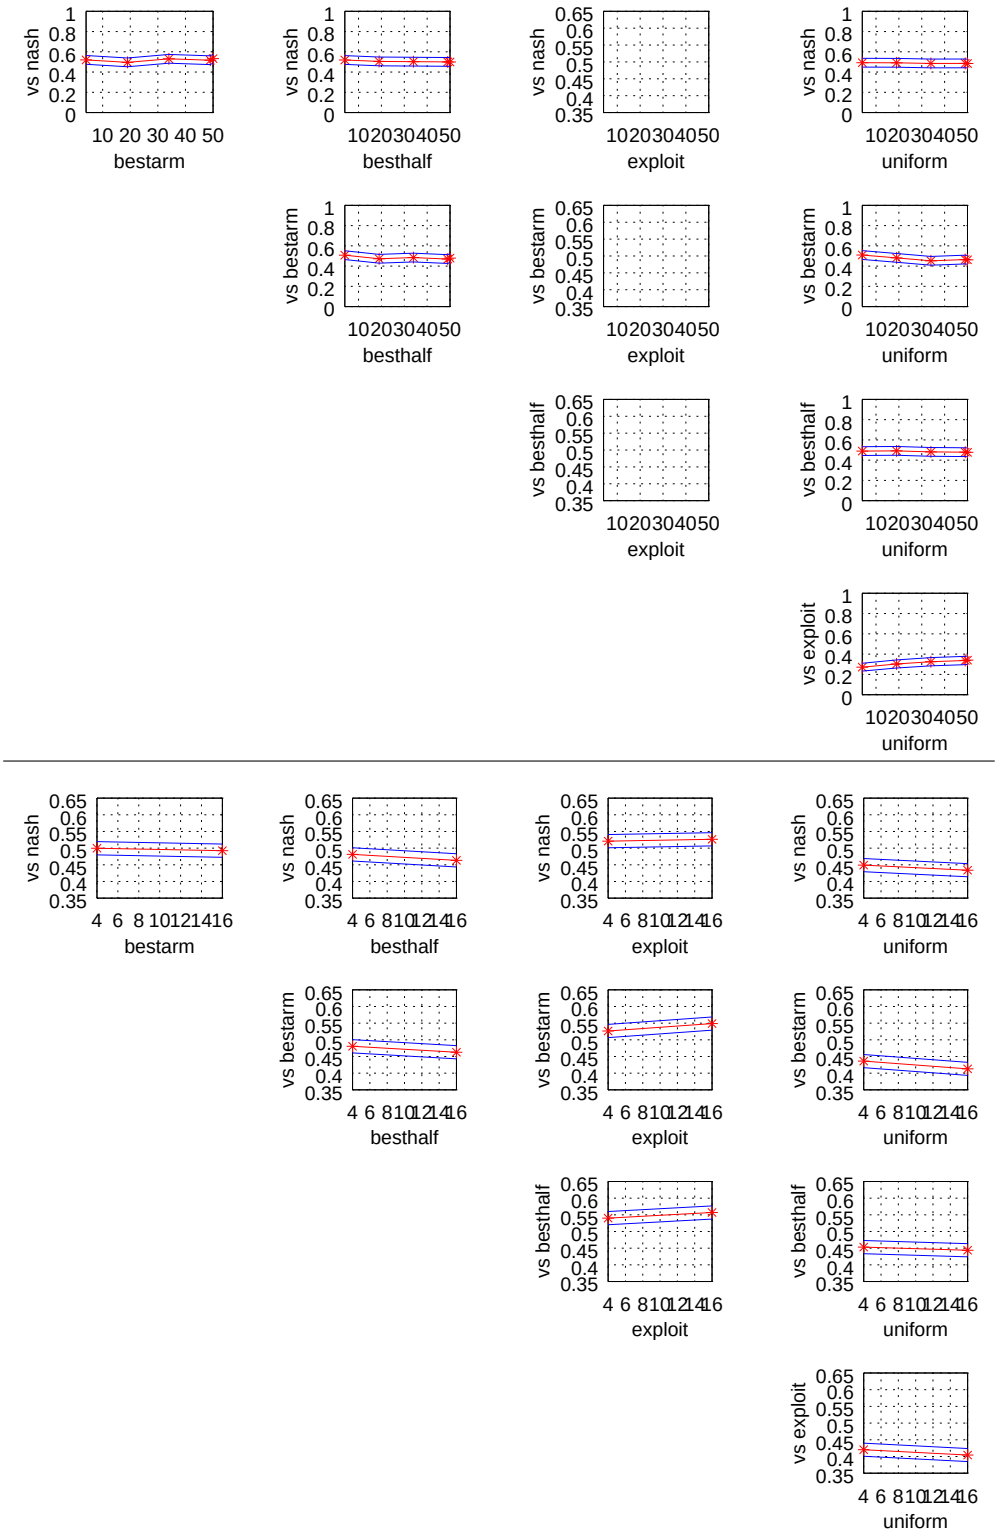

Fig. 10: Top: Chess, portfolio of random seeds: performance of various portfolio algorithms (see text) against others, depending on the number of options in the solver (x-axis). Bottom: Go, portfolio of GnuGo variants: performance of various portfolio algorithms (see text) against others, depending on the number of options in the solver (x-axis).

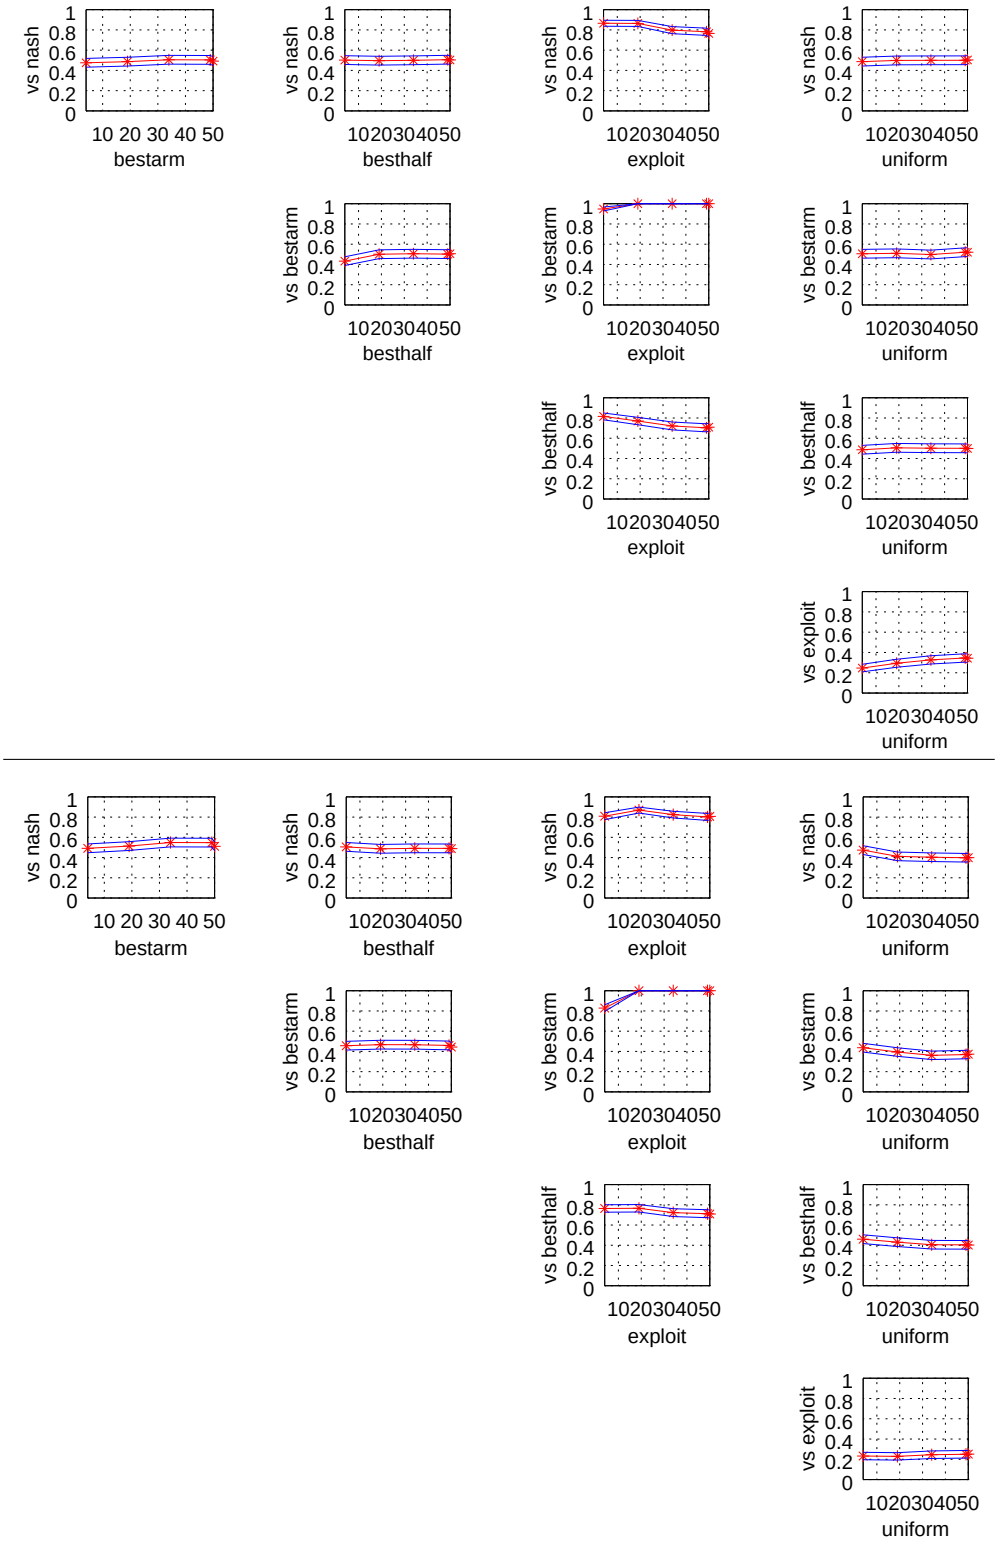

Fig. 11: Top: Havannah, portfolio of random seeds: performance of various portfolio algorithms (see text) against others, depending on the number of options in the solver (x-axis). Bottom: Batoo, portfolio of random seeds: performance of various portfolio algorithms (see text) against others, depending on the number of options in the solver (x-axis).